

ARI Research Note 96-51

Natural Language Access to Intelligent Systems

George A. Miller
Princeton University

Research and Advanced Concepts Office
Michael Drillings, Acting Director

March 1996

19960828 032



United States Army
Research Institute for the Behavioral and Social Sciences

Approved for public release; distribution is unlimited.

U.S. ARMY RESEARCH INSTITUTE FOR THE BEHAVIORAL AND SOCIAL SCIENCES

**A Field Operating Agency Under the Jurisdiction
of the Deputy Chief of Staff for Personnel**

EDGAR M. JOHNSON
Director

Research accomplished under contract
for the Department of the Army

Princeton University

Technical review by

Joseph Psotka

DTIC QUALITY INSPECTED 3

NOTICES

DISTRIBUTION: This report has been cleared for release to the Defense Technical Information Center (DTIC) to comply with regulatory requirements. It has been given no primary distribution other than to DTIC and will be available only through DTIC or the National Technical Information Service (NTIS).

FINAL DISPOSITION: This report may be destroyed when it is no longer needed. Please do not return it to the U.S. Army Research Institute for the Behavioral and Social Sciences.

NOTE: The views, opinions, and findings in this report are those of the author(s) and should not be construed as an official Department of the Army position, policy, or decision, unless so designated by other authorized documents.

REPORT DOCUMENTATION PAGE

1. REPORT DATE 1996, March		2. REPORT TYPE Final		3. DATES COVERED (from... to) July 1986-June 1989	
4. TITLE AND SUBTITLE Natural Language Access to Intelligent Systems				5a. CONTRACT OR GRANT NUMBER MDA903-86-K-0242	
				5b. PROGRAM ELEMENT NUMBER 0601102A	
6. AUTHOR(S) George A. Miller (Princeton University)				5c. PROJECT NUMBER B74F	
				5d. TASK NUMBER 8522	
				5e. WORK UNIT NUMBER C71	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Department of Psychology Princeton University Princeton, NJ 08544-1010				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) U.S. Army Research Institute for the Behavioral and Social Sciences ATTN: PERI-BR 5001 Eisenhower Avenue Alexandria, VA 22333-5600				10. MONITOR ACRONYM ARI	
				11. MONITOR REPORT NUMBER Research Note 96-51	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution is unlimited.					
13. SUPPLEMENTARY NOTES COR: Judith Orasanu					
14. ABSTRACT (<i>Maximum 200 words</i>): Work under this contract had two components, both aimed at facilitating natural language access to intelligent systems. One aspect was concerned with increasing the vocabularies of personnel who use intelligent systems, and the other was an attempt to increase the vocabulary that systems can process intelligently.					
15. SUBJECT TERMS Vocabulary Semantic memory Learning from context Interactive videodisc					
SECURITY CLASSIFICATION OF			19. LIMITATION OF ABSTRACT Unlimited	20. NUMBER OF PAGES 13	21. RESPONSIBLE PERSON (Name and Telephone Number)
16. REPORT Unclassified	17. ABSTRACT Unclassified	18. THIS PAGE Unclassified			

NATURAL LANGUAGE ACCESS TO INTELLIGENT SYSTEMS

Army Research Institute, Contract MDA 903-86-K-0242 (7/86 - 6/89)
Behavioral and Social Sciences

Final Technical Report
31 December 1989

Submitted by:

George A. Miller, Principal Investigator
Cognitive Science Laboratory
Department of Psychology
Princeton University
Princeton, New Jersey 08544

Contents

Vocabulary Growth in Human Language Users	2
Videodisc tutoring	2
Picturability	3
Sentential contexts	3
Lexical tutoring	4
WordNet: A Lexical Database	5
The Organization of Lexical Knowledge	5
Nouns	7
Adjectives	7
Verbs	8
Lexpert	8
Technical Reports	9
Publications	10

NATURAL LANGUAGE ACCESS TO INTELLIGENT SYSTEMS

Work under this contract had two components, both aimed at facilitating natural language access to intelligent systems. One aspect was concerned with increasing the vocabularies of personnel who use intelligent systems, the other was an attempt to increase the vocabulary that computer systems can process intelligently.

Work under this contract was scheduled to terminate as of 30 June 1989. In order to bring the research to an orderly conclusion and to prepare a final report on the work, however, a no-cost extension for six months was requested and granted.

Vocabulary Growth in Human Language Users

Prior research had established that learners who are unfamiliar with a word are at risk if they try to use that word solely on the basis of a dictionary definition. Vocabulary growth occurs primarily through observations of how words are used by others: used in speech throughout life, and used in writing after achieving literacy. For that reason, emphasis in the present research was placed on understanding and facilitating the cognitive processes whereby words are learned from context.

(1) *Videodisc tutoring.*

In discussions of vocabulary acquisition, "context" has a narrow and a broad sense. The narrow sense includes only linguistic contexts—the words uttered before and after the word in question. The broad sense also includes linguistic contexts, but in addition includes any non-linguistic information that a learner may have about the situation and the participants and the participants' communicative intentions. It is arguable that children do not acquire language when nothing more than purely linguistic context is available to them. Be that as it may, it is certainly true that children learn their first words with very rich environmental support.

One goal of the research, therefore, was to provide a broad context and to evaluate its importance for word learning. We decided to enrich the contextual support for learners by embedding the words to be learned in a narrative, rather than in isolated sentences. And we decided to provide visual context as well as linguistic context—for that purpose we used interactive videodisc. The general plan was to ask learners to read a text that described an episode from a movie that they had just seen. Included in the text were certain words, specially marked, that the reader was expected to learn. When one of those target words was encountered, various kinds of information about its meaning were available. In order to assess the value of different kinds of information about the target words, a pretest was given before the movie episode was shown and a posttest assessed comprehension of the story and memory of the target words.

The major experimental variable was the kind of help available to learners. Three kinds of assistance were offered: (a) dictionary definitions; (b) pictures; or (c) other sentences using the word in the same sense. When these three kinds of help were offered freely the learners asked to see definitions; they seldom asked for illustrative sentences.

In order to make the desired comparison, it was necessary to test separate groups, where each group received one pattern of assistance. A pattern consisted of one or more of the three kinds of help that were available. It was found that dictionary definitions produced little learning; the most effective help was another sentence suitable to the same story and using the target word in the same sense. Seeing a further sentence seemed to challenge the learners to think about it and to form connections between the word and things they already knew. Combining two or all three kinds of help was not effective—learners seemed to choose the one kind they liked best (i.e., definitions) and ignored the others. These results were obtained both for children and for college freshmen.

The disappointing results obtained when pictures were used was unexpected in view of the importance traditionally attached to ostensive definitions. In part, this poor result was attributable to our use of commercially recorded motion pictures, where we had little freedom to generate the pictures that would have been optimal for our purposes. We considered the possibility of developing our own instructional pictures, but decided not to pursue that line. Our subjects already knew names for most concrete picturable objects; the words we were trying to teach were more abstract, low-imagery words. We needed to know more about the picturability of relatively abstract concepts.

(2) Picturability.

In order to explore the usefulness of pictures for vocabulary instruction, a series of studies was conducted on picturability. The experiments compared the usefulness of verbal and visual contexts in eliciting specific lexical targets from subjects. The results showed that linguistic contexts are much better at activating specific lexical targets than pictures are. Not only do definitions elicit the correct target word more frequently, but also do so more rapidly. Pictures activate a wider range of associations, but do not direct the learner to the specific term in question. Only when the target is extremely concrete do subjects feel that a picture is a better source of information about the target than is a definition. Since most of the relatively common words that our subjects did not already know were more abstract, improving the quality of the pictures did not seem to be a worthwhile strategy to pursue.

(3) Sentential contexts.

Subsequent research focussed, therefore, on linguistic contexts. Most previous work on the relation of words to their linguistic contexts had used a single context and investigated the constraints that that context imposed on the various words that could be substituted into it. This seemed to reverse the situation that a learner normally faces: a single word is at issue, and it is observed to occur in a variety of different contexts.

In order to simulate a normal learning situation more closely, therefore, a new technique was developed, called "the method of sorting linguistic contexts." In its simplest form, (a) a pair of words is chosen, (b) representative samples of sentences using that word are collected, (c) the target words are deleted, (d) the resulting contexts are shuffled together, and (e) subjects are challenged to sort them out. The results permit estimations of the probability of a correct recognition and of a false alarm; from those estimated probabilities a measure d' of the discriminability of the contexts is calculated according to signal detection theory.

An extensive series of studies demonstrated the robustness of this technique. The number of words to be discriminated can be increased beyond two; it makes little absolute difference whether the subjects know in advance what the target words are; sentences can be collected either from printed corpora or composed on request by college students. It was found that different senses of the same word (e.g., *board* as a committee vs. *board* as a piece of wood) can be accurately discriminated—this result confirms the general observation that context serves to disambiguate polysemous words. Response biases (e.g., when a pair like *hotel/building* is tested, there is a strong favoritism for sorting contexts of *hotel* into the pile of *building* contexts, but not vice versa) were found to be attributable to familiarity as indicated by relative frequencies of use—their effect was minimized by the use of signal detection theory. And words that seem intuitively to be related in meaning generally occur in contexts that are difficult to discriminate—the method of sorting provides a way to explore people's intuitive judgments of semantic similarity.

The method of sorting was used to study antonyms. Although in principle any noun that can be modified by one adjective can also be modified by its antonym, in fact it was found that the contexts of antonymous adjectives are highly discriminable. This finding raised a question as to how antonyms come to be so closely associated (e.g., in word association experiments), and an analysis of the one-million-word Brown Corpus revealed that antonyms co-occur in sentences far more frequently than chance would predict. The finding also confirmed the observation that contextual discriminability varies inversely with semantic similarity.

The method of sorting was used to study pairs of nouns that differed widely in their judged semantic similarity. Previous attempts to explore the relation between semantic similarity and contextual similarity had used a measure of contextual similarity based on the similarity of the words that co-occur with the target words, but the relation was found only for very similar target words. With the method of sorting, which provides a measure of contextual similarity based on substitutability rather than co-occurrences, a linear function was found over the entire range of similarities. These results led to the formulation of a general theory of contextual representations that will be tested and developed in further research.

(4) *Lexical Tutoring.*

Since it had been found that dictionary definitions are more confusing than helpful for a learner and that pictures are of little help for learning the meanings of generic or abstract words, our efforts were re-directed toward optimizing the conditions under which example usages of a word could be used for learning from context. Since video-disc technology is expensive and not universally available, dispensing with it had certain practical advantages. A more difficult modification involved getting rid of the experimenter.

In our initial instructional interventions, an experimenter had always been available to record the subjects' responses and to offer advice and encouragement. A series of studies was conducted to develop a technique that did not rely on the presence of another person. From video recording of subjects it was found that learners who were given a

definition quickly returned to reading the story; the meaning of the word was secondary to getting on with the narrative. Learners who were given illustrative sentences, however, were much more active; they spent more time thinking about the word and seemed to be engaging in a discovery process. Learners who saw only sentences took approximately 20 seconds longer than did learners who saw only definitions.

A procedure was tested, therefore, that involved giving alternative sentences and pausing for 20 seconds while the reader thought about the meaning of the word. This tactic appeared to be successful without the intervention of a coach or experimenter to elicit vocal responses involving the word. It was found to be most effective for teaching nouns and least effective for teaching adverbs. The delay method is presently being incorporated into a handheld teaching device that will be tested in further experiments. If the results are encouraging, it is hoped that the handheld device can be made commercially available.

WordNet: A Lexical Database

A second component of the work under this contract concerned the development of a lexical database that could be used in processing natural language by computers. It was initially proposed to provide procedural definitions written in LISP code that could be incorporated into larger language-understanding programs. When this goal proved too ambitious for practical realization, it was revised. As a first step toward the original goal, work was directed toward compiling an on-line lexical database, called WordNet, that is organized according to psychological theories of semantic memory.

How WordNet could be used to achieve some of the goals originally envisioned for this project can be illustrated by an example. Suppose that a computer encountered the following two sentences: *It is an attack fighter. The wings are aluminum.* A computer could easily parse these sentences, but it would have no reason to see any relation between them. Given access to lexical information, however, the computer should be able to recognize that one sense of *fighter* is *airplane*, and that an *airplane* has *wings*, thus discovering the connection that is immediately obvious to a human reader. The goal of WordNet is to provide that kind of lexical information in a form usable by computers.

Work on WordNet was also supported by contract N00014-86-K-0492 with the Office of Naval Research, and is presently continuing under that contract.

(1) *The Organization of Lexical Information.*

A word is an association of a meaning with an utterance (or inscription) that plays a grammatical role—that much is irreducible. The basic plan of WordNet, therefore, can be visualized as a vocabulary matrix formed by N words \times M meanings, where an entry in a cell means that the word in that column can be used to express the meaning in that row of the matrix. Any lexical database—printed dictionaries, computerized lexical databases, human semantic memory—can be represented in this fashion.

A critical problem facing anyone who would build on this representation is how the meanings (or senses) of words should be characterized. In a printed dictionary, the

meanings are given in the form of short glosses, supplemented by illustrative phrases or sentences. In a person's memory, the representation of lexical meanings is still a matter for debate, although some theorists have claimed to write LISP programs that characterize what a person knows about words. For WordNet, a compromise was reached: a meaning is represented by the set of synonyms that can be used to express it. This solution is transparent to human users; it is of less value for computer processing. Still, the synonym sets (synsets) provide convenient entries between which semantic relations can be represented by labelled pointers.

Many different semantic relations have been discussed in the literature, but a limited set of semantic relations was used to organize WordNet. One criterion for inclusion in the set was that the relation should be familiar to lay persons, i.e., that it should not be necessary to have special linguistic training in order to understand the semantic relations that were being used. Another criterion was that the relation should be reciprocal, i.e., if the relation R holds between word W_1 and word W_2 , then a converse relation R' must hold between W_2 and W_1 .

Synonymy: The semantic relation of synonymy is built into WordNet by the synsets that form the basic building blocks of the database. The relation of synonymy is its own converse: if W_1 is a synonym of W_2 , then W_2 is a synonym of W_1 . Synonymy holds between nouns, adjectives, and verbs. A special case of synonymy is required for adjectives, for reasons that will be explained below.

Hyponymy/Superordination: If N_1 is a kind of N_2 , then N_1 is said to be a hyponym of N_2 . The relation of superordination (hypernymy) is the converse of hyponymy: if N_1 is a hyponym of N_2 , then N_2 is a superordinate of N_1 . Although this semantic relation holds only for nouns, it is transitive, which makes it important for inferring logical consequences.

Meronymy/Holonymy: If N_1 is a part of N_2 , then N_1 is said to be a meronym of N_2 . The relation of holonymy is the converse of meronymy: if N_1 is a meronym of N_2 , then N_2 is a holonym of N_1 . This semantic relation holds only for nouns. There are a variety of part/whole relations that can be expressed by meronymy, and transitivity holds only within a single kind (e.g., a *branch* is a part of a *tree* and a *tree* is a part of a *forest*, but a *branch* is not a part of a *forest*; one relation is between a component and a whole, whereas the other is between a member and a collection).

Antonymy: The semantic relation of antonymy is its own converse: if A_1 is an antonym of A_2 , then A_2 is also an antonym of A_1 . Although some nouns and verbs have antonyms, antonymy is the basic organizing relation for adjectives.

Troponymy: If V_1 is to V_2 in some manner (e.g., to *weigh* is to *measure* in some manner), then V_1 is a troponym of V_2 . The semantic relation of troponymy resembles, and is often mistaken for, the relation of hyponymy. The difference is that troponymy holds between verbs, whereas hyponymy holds between nouns. For example, *weight* is a kind of *measure*, and *weighing* can be said to be a kind of *measuring*. But one would not say *measure* is a part of *weight*, although most people would understand that *measuring* is a part of *weighing*. Since nouns and verbs behave differently in this respect, a special relation of troponymy has been introduced for the verbs. No technical term has yet been

adopted for the converse of troponymy.

Entailment: Entailment (strict implication) is defined for propositions—a proposition P is said to entail a proposition Q if and only if there is no conceivable state of affairs that could make P true and Q false. In WordNet, entailment is generalized to denote the semantic relation between two verbs V_1 and V_2 that holds when the statement *Someone V_1 -s* entails the statement *Someone V_2 -s*.¹ No technical term has yet been adopted for the converse of entailment.

These are the semantic relations that have been used to organize WordNet. Since it was found that different relations are needed to characterize the organization of different parts of speech, nouns, verbs, and adjectives will be discussed separately.¹

(2) Nouns.

There are approximately 32,000 nouns in WordNet, organized into some 25,000 synsets. Although all of the familiar nouns are entered (in several senses), words are still being added as the work continues. The nouns are topically organized into 25 separate lexical files: noun.Tops, noun.act, noun.animal, noun.artifact, noun.attribute, noun.body, noun.cognition, noun.communication, noun.event, noun.feeling, noun.food, noun.group, noun.location, noun.motive, noun.object, noun.person, noun.phenomenon, noun.plant, noun.possession, noun.process, noun.quantity, noun.relation, noun.shape, noun.social, noun.state, noun.substance, and noun.time.

Each file is organized by hyponymy into an extended and multiply branching tree, a hierarchical organization that is familiar from taxonomies and from expert systems. In WordNet it is not only possible to find superordinates (to find that *tree* is a species of *plant*), but also to find lists of hyponyms (to enter with *tree* and find a list of kinds of trees), which is very difficult with most printed dictionaries.

In some of the files—body parts and human artifacts being the most extreme examples—the problem arises of relating hyponymy and meronymy. Since the hyponymic tree allows logical inheritance, in order to find all the parts of something X, it is necessary to search up the tree to find all the parts that X inherits from its superordinates. Thus, attaching the meronyms to the hierarchy at the right node is often a challenging task. In this respect, WordNet is not a faithful model of human semantic memory, because people organize parts by imagery—and imagery is not available in WordNet.

(3) Adjectives.

WordNet contains well over 13,000 adjectives organized into some 10,000 synsets. All adjectives are together in a single large file.

¹ Adverbs were not included in WordNet because the vast majority of them are simply adjectives with an *-ly* suffix. The minor parts of speech were not included in WordNet because they provide important clues to the syntactic structure of sentences and so will play an important role in the parser. The irregular adverbs and minor parts of speech are limited in number and can easily be added in a manner compatible with any language-processing system that will incorporate WordNet.

When we first turned our attention to adjectives we were surprised to discover that the hierarchical scheme so successful in organizing the noun lexicon was useless for adjectives. The problem is easily illustrated.

Antonymy is the basic semantic relation organizing adjectives. Suppose we have two synsets: { *wet, moist* } and { *dry, arid* }. The simple way to represent antonymy would be to introduce a labelled pointer between these synsets. But while that would work well for *wet/dry*, most people feel that *wet/arid* and *moist/dry* are much weaker antonyms, and *moist/arid* are hardly antonyms at all. The method we had been using would not work for adjectives. By a painfully roundabout way we had discovered what every lexicographer knows, namely, that antonymy is a relation between words, not a relation between meanings, senses, or concepts.

Adjectives express values of attributes, and most attributes are bipolar. Thus, there is a cluster of adjectives around each opposing pole of the attribute. Usually, a pair of adjectives (often more than one pair) will be chosen for a special associative bonding (presumably learned from their frequent co-occurrence in the same sentences), and the others will cluster around them and express more specific or subtler values of the attribute. In order to represent this structure in WordNet, we distinguished between direct antonyms (e.g., *wet/dry*) and indirect antonyms (e.g., *wet/arid, moist/dry*). Indirect (or conceptual) antonymy is mediated by two semantic relations, one of similarity (a loose form of synonymy) to an adjective that does have an antonym, and another of direct antonymy.

(4) *Verbs.*

WordNet contains 4,500 verbs organized into 3,300 synsets. The verbs are organized into fourteen separate lexical files: *verb.body*, *verb.change*, *verb.cognition*, *verb.communication*, *verb.competition*, *verb.consumption*, *verb.contact*, *verb.creation*, *verb.motion*, *verb.perception*, *verb.possession*, *verb.psych*, *verb.social*, and *verb.weather*.

Verbs have a complex semantic organization that has been difficult to deal with. Previous work had concentrated on the syntactic properties of verbs (e.g., transitive vs. intransitive) or on their semantic decomposition (e.g., *kill* is decomposed into *cause* and *die*), and far less attention had been paid to the semantic relations between verbs. Special skill is needed to deal with verbs that seem to be closely related in meaning, yet have very different syntactic properties.

One feature of the verb files is the provision of sentence frames (e.g., *Someone — something to someone*). With each verb is associated a frame or set of frames in which the verb can be used. This mechanism provides a fuller account of the transitive/intransitive distinction than is found in most desk dictionaries.

(5) *Lexpert.*

The software involved in creating an interface between a user and the lexical database is referred to (loosely) as Lexpert. Several kinds of software have been developed: (a) programs to assist lexicographers in entering words into the lexical files, and in editing those files; (b) the "grinder" programs that convert lexical files into the actual database; (c) programs to search the database and retrieve the requested information; (d)

programs to generate the display that the user sees and reacts to.

The grinder constructs a masterlist containing all of the words in any of WordNet's lexical files, along with the addresses of any information that WordNet contains about each word. Polysemous words are entered once for each meaning, i.e., once for each synset the word belongs to. For each word, the masterlist also contains information about the word's part of speech, the word's familiarity, and all of the semantic relations that the word enters into. Some of the semantic relations that are listed have to be invented by the grinder. That is to say, lexicographers enter semantic relations in the lexical files in only one direction (e.g., only a superordinate-pointer from *maple* to *tree*) and the converse pointer (the hyponymic-pointer in this example) is added to the database automatically by the grinder.

The basic task of the search code is to look in the masterlist for the requested word and to tell the display what kind of information is available. When the user makes a choice (e.g., from a menu) the search code goes to the appropriate addresses listed on the masterlist and retrieves the requested information. A problem with this scheme that has not yet been solved, but is being dealt with as this contract ends, arises from the fact that not all forms of given word are listed. That is to say, it should be possible for a user to put the cursor on *went* and retrieve the information that WordNet has available about the verb *go*. The software needed to do this, dubbed Murphy by its creator, will be an important component of any practical applications of WordNet.

Finally, the display interface is peculiar to the particular computer that a user has access to. The system has been developed on Sun workstations, using the windowing routines available in SunTools. But alternative interfaces are being written for Macintosh and MS-DOS computers. It is estimated that approximately one more year of work will be required to bring WordNet and its associated software to the point that it can be made available to any qualified research workers who have a use for it.

Additional detailed information not contained in this executive summary can be found in the Quarterly Progress Reports and in the reports and publications listed below.

Technical Reports

The following is a list of Technical Reports from the Princeton Cognitive Science Laboratory that were supported, in whole or in part, under this contract:

- Beckwith, R., Fellbaum, C., Gross, D., & Miller, G. A. (1990). *WordNet: A Lexical Database Organized on Psycholinguistic Principles*. CSL Report 42. 12 pp.
- Bienkowski, M. A. (1987). *Tools for Lexicon Construction*. CSL Report 10. 18 pp.
- Charles, W. G., & Miller, G. A. (1988). *Contexts of antonymous adjectives*. CSL Report 32. 24 pp.
- Collier, G. H., & Fellbaum, C. (1988). *Exploring the Verb Lexicon with the Sensus Electronic Thesaurus*. CSL Report 30. 10 pp.
- Gildea, P. M., Miller, G. A., & Wurtenberg, C. L. (1988). *Contextual Enrichment by Videodisc: A First Report*. CSL Report 12. 22 pp.

- Gross, D., Fischer, U., & Miller, G. A. (1988). *Antonymy and Representation of Adjectival Meanings*. CSL Report 13. 15 pp.
- Hanson, C. (1988). *Eliciting Lexical Targets: A Comparison of Definitions and Depictions*. CSL Report 19. 16 pp.
- Hanson, C., and Fellbaum, C. (1988). *Can You Picture It? A Look at Idioms*. CSL Report 26. 22 pp.
- Miller, G. A., Fellbaum, C., Kegl, J., & Miller, K. (1988). *WordNet: An Electronic Lexical Reference System Based on Theories of Lexical Memory*. CSL Report 11. 18 pp.
- Teibel, D. A. (1988). *WordNet User's Guide*. CSL Report 34. 15 pp.
- Teibel, D. A. (1988). *A Multilayered Approach to Constructing a Representation of the English Lexicon*. CSL Report 35. 22 pp.

Publications

The following is a list of journal publications that were supported, in whole or in part, under this contract:

- Charles, W. G. (1988). The categorization of sentential contexts. *Journal of Psycholinguistic Research*, 17, 403-411.
- Charles, W. G., & Miller, G. A. (1989). Contexts of antonymous adjectives. *Applied Psycholinguistics*, 10, 357-375.
- Gildea, P. M., Miller, G. A., & Wurtenberg, C. L. (1990). Contextual enrichment by videodisc. In Nix, D., & Spiro, R. J. (eds.) *Cognition, Education and Multimedia*. Hillsdale, NJ: Erlbaum. Pp. 1-29.
- Gross, D., Fischer, U., & Miller, G. A. (1989). The organization of adjectival meanings. *Journal of Memory and Language*, 28, 92-106.
- Miller, G. A., & Charles, W. G. (submitted). Contextual correlates of semantic similarity. *Language and Cognitive Processes*.
- Miller, G. A., Fellbaum, C., Kegl, J., & Miller, K. (1988). WordNet: An electronic lexical reference system based on theories of lexical memory. *Revue quebecoise de linguistique*, 17, 181-213.